arXiv:1608.08126v1 [stat.ME] 29 Aug 2016

Simultaneous penalized M-estimation of covariance matrices using geodesically convex optimization

Esa Ollila[a,*], Ilya Soloveychik[b], David E. Tyler[c,1], Ami Wiesel[b,2]

[a] *Aalto University, Finland*
[b] *The Hebrew University of Jerusalem, Israel*
[c] *Rutgers – The State University of New Jersey, USA*

Abstract

A common assumption when sampling p-dimensional observations from K distinct group is the equality of the covariance matrices. In this paper, we propose two penalized M-estimation approaches for the estimation of the covariance or scatter matrices under the broader assumption that they may simply be close to each other, and hence roughly deviate from some positive definite "center". The first approach begins by generating a pooled M-estimator of scatter based on all the data, followed by a penalised M-estimator of scatter for each group, with the penalty term chosen so that the individual scatter matrices are shrunk towards the pooled scatter matrix. In the second approach, we minimize the sum of the individual group M-estimation cost functions together with an additive joint penalty term which enforces some similarity between the individual scatter estimators, i.e. shrinkage towards a mutual center. In both approaches, we utilize the concept of geodesic convexity to prove the existence and uniqueness of the penalized solution under general conditions. We consider three specific penalty functions based on the Euclidean, the Riemannian, and the Kullback-Leibler distances. In the second approach, the distance based penalties are shown to lead to estimators of the mutual center that are related to the arithmetic, the Riemannian and the harmonic means of positive definite matrices, respectively. A penalty based on an ellipticity measure is also considered which is particularly useful for shape matrix estimators. Fixed point equations are derived for each penalty function and the benefits of the estimators are illustrated in regularized discriminant analysis problem.

Keywords: discriminant analysis, geodesic convexity, M-estimators of scatter matrix, shrinkage, regularization

*Corresponding author
Email address: esa.ollila@aalto.fi (Esa Ollila)

[1]Research partially supported by the National Science Foundation Grant No. DMS-1407751. Any Opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect those of the National Science Foundation.

[2]The author gratefully acknowledge Israel Science Foundation grant 1339/15

1. Introduction

Many multivariate statistical applications require the simultaneous estimation of the covariance matrices $\mathbf{\Sigma}_1, \ldots, \mathbf{\Sigma}_K$ of a set of multivariate measurements on K distinct groups. Often the sample sizes $n_k, k = 1, \ldots, K$, of each group are small relative to the dimension p, which makes estimating the individual covariance matrices a challenge. Quite often, though, based on the physical properties of the underlying measured phenomena or experience with similar datasets, one may postulate the existence of common features or similarities among the estimated covariance matrices. This prior knowledge can be incorporated into the estimation problem by either modeling the covariance matrices as having some common structure or by pooling the data from the K groups.

In this paper we focus on data pooling techniques via regularization. The use of pooling and regularization methods assume the distinct covariance matrices share some common features, without necessarily modeling the common features. A prominent example of this approach is Friedman's regularized discriminant analysis [1]. A similar approach to estimating precision matrices, i.e. inverse covariance matrices, was treated in [2]. The goal of Friedman's regularized discriminant analysis approach is to strike a balance between quadratic and linear discriminant analysis (QDA/LDA) in the under-sampled scenario via shrinkage regularization. In [1] it was illustrated that it is often beneficial to shrink the class Sample Covariance Matrices (SCM) towards the pooled SCM.

The methods proposed in these works were developed under the assumption of sampling from multivariate normal distributions. Consequently, they tend to depend on variants of the SCM estimator and are not resistant to outliers nor robust against heavier tailed distributions. From this perspective, taking into account the non-Gaussianity of measurements in many real world applications, the statistical community has become increasingly aware of the advantage of more robust and resistant multivariate methods. This, in particular, led to development of the family of the M-estimators of multivariate scatter [3, 4, 5], as well as families of high-breakdown point scatter estimators such as the MVE and MCD estimator [6], the S-estimators [7], and the MM-estimators [8], among others. There appears, though, to be little work on robustness in the context of joint covariance estimation and its application to regularized discriminant analysis and other problems. The intent of this paper is to address this issue. We focus on M-estimation methods, which unlike the high breakdown point methods, are readily amenable to the sparse data setting and regularization.

Our aim is to propose robust versions of the SCM based shrinkage covariance estimators proposed in [1] for regularized discriminant analysis (RDA) in the sparse data setting. The approach used in [1] is based on taking a convex combination of the individual SCM and the pooled SCM. Such an approach, though, does not directly generalize when using M-estimators of scatter, since the M-estimators are not defined when the data within a group is sparse. Rather, in our approach we apply penalization to M-estimation loss functions. When

using such loss functions which correspond to bounded influence M-estimators of scatter, though, one encounters a non-convex optimization problem in Euclidean space. Here, the concept of geodesic convexity (g-convexity) plays a crucial role, which basically means switching to a different metric over the set of positive definite matrices, for which the loss function is then convex in this metric. The use of g-convexity in covariance estimation was introduced in [9] and has subsequently been utilized in related works, e.g. [10, 11, 12, 13]. See [14] for a nice overview of usage of g-convexity in covariance matrix estimation problems. Introducing additive g-convex penalty terms to the loss functions, keeps the optimization problem g-convex.

Two penalized M-estimation approaches are introduced for the problem of joint estimation of group covariance matrices. The first approach begins by defining a pooled M-estimator of scatter based on all the data, followed by a penalized M-estimator of scatter for each group, with the penalty term chosen so that the individual scatter matrices are shrunk towards the pooled scatter matrix. In the second approach, we minimize the sum of the individual group M-estimation loss functions together with an additive joint penalty term which enforces some similarity between the individual scatter estimators, i.e. shrinkage towards a mutual center. Hence, in the second approach, the individual covariance matrices and their mutual center are estimated simultaneously. In both approaches, we consider three g-convex penalty functions based on the Euclidean, the Riemannian, and the information theoretic (Kullback-Leibler) distances. In the second approach, these penalties are shown to lead to estimators of the mutual center that are related to the arithmetic, the Riemannian and the harmonic means of positive definite matrices, respectively. We also consider a penalty based on an ellipticity measure for positive definite matrices, which shrinks the individual estimators towards a common shape matrix rather than a common scatter matrix.

The rest of the paper is organized as follows. Section 2 introduces our penalized M-estimation approaches for estimating the unknown K scatter matrices $\{\mathbf{\Sigma}_k\}_{k=1}^K$ and their joint center $\mathbf{\Sigma}$. Examples of g-convex loss functions, including the Gaussian, Huber's and Tyler's loss functions, are given. Section 3 provides a brief introduction to g-convex functions of positive definite symmetric (PDS) matrices. In Section 4 examples of g-convex penalty/distance functions are given. In addition, we show that the KL-distance and the ellipticity distance are g-convex, and when used for defining a center for a given $\{\mathbf{\Sigma}_k\}_{k=1}^K$ yield weighted harmonic means of positive definite matrices. In Section 5 we derive general conditions for uniqueness of the solution as well as derive fixed point algorithms for their computation. Section 5.1 considers existence and uniqueness conditions separately for Tyler's loss function. Section 6 describes a cross validation procedure for penalty parameter selection. In Section 7 we illustrate the application of the proposed scatter matrix estimators to regularized discriminant analysis and illustrate the performance of RDA rules via a small simulation study and a data example. Section 8 concludes the paper. Proofs are given in the Appendix.

Notation: Let $\mathcal{S}(p)$ be the open cone of positive definite $p \times p$ symmetric ma-

trices, and let $\mathbf{I}$ be the identity matrix of proper dimension. On $\mathcal{S}(p)$, we denote the Frobenius norm by $\|\cdot\|_{\mathrm{F}}$, the spectral norm by $\|\cdot\|_2$, and the determinant by $|\cdot|$.

2. Problem Formulation

2.1. General Setting

The multivariate M-estimators were introduced in [4] as generalizations of the maximum likelihood estimators for an elliptically symmetric multivariate distribution. An absolutely continuous random vector $\mathbf{x} \in \mathbb{R}^p$ is said to have a real elliptically symmetric (RES) distribution with center of symmetry $\boldsymbol{\mu}$ and scatter matrix parameter $\boldsymbol{\Sigma} \in \mathcal{S}(p)$, if it has a density of the form

$$f(\mathbf{x}|\boldsymbol{\Sigma}) = C_{p,g}|\boldsymbol{\Sigma}|^{-1/2} g\{(\mathbf{x}-\boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1}(\mathbf{x}-\boldsymbol{\mu})\}, \tag{1}$$

where $C_{p,g}$ denotes the normalizing constant, and $g : \mathbb{R}^+ \to \mathbb{R}^+$ is viewed as a density generator. Here, $\mathbb{R}^+ = \{x \in \mathbb{R} | x \geq 0\}$. For simplicity, we state $\mathbf{x} \sim \mathcal{E}_p(\boldsymbol{\mu}, \boldsymbol{\Sigma}, g)$. The function g determines the radial distribution of the elliptical population and hence the degree of its "heavy-tailedness". The scatter matrix $\boldsymbol{\Sigma}$ is proportional to the covariance matrix whenever the second moments exist, and serves as a generalization of the covariance matrix when the second moments do not exist. There is a extensive literature on the properties of elliptical distributions. The elliptical family includes many widely used multivariate distributions such as as Gaussian, compound Gaussian, K-distributions, among many others. For a thorough treatment of elliptical distributions and their generalizations see e.g., [15, 16].

Consider samples from K distinct groups of p-dimensional measurements,

$$\mathbf{x}_{11}, \ldots, \mathbf{x}_{1n_1}, \quad \ldots, \quad \mathbf{x}_{K1}, \ldots, \mathbf{x}_{Kn_K}, \tag{2}$$

with group $\mathbf{X}_k = \{\mathbf{x}_{k1}, \ldots, \mathbf{x}_{kn_k}\}$ have sample size n_k, $k = 1, \ldots, K$. Let

$$N = \sum_{i=1}^{K} n_k \quad \text{and} \quad \pi_k = \frac{n_k}{N}, \text{ for } k = 1, \ldots, K \tag{3}$$

denote the total sample size and the relative sample sizes of each of the K groups, respectively. The measurements are assumed to be mutually independent and within each group they are assumed to be identically distributed.

In our development, we first presume the measurements within the different groups follow elliptical distributions with known centers of symmetry, which we take without loss of generality to be $\boldsymbol{\mu}_k = \mathbf{0}$ for $k = 1, \ldots, K$. The assumption of having known centers is to be discussed later. Hence, we assume the random sample of the measurements for the kth group comes from an $\mathcal{E}_p(\mathbf{0}, \boldsymbol{\Sigma}_k, g_k)$ distribution, $k = 1, \ldots, K$, with possibly different scatter matrices $\boldsymbol{\Sigma}_k$. The

negative log-likelihood for this scenario, ignoring the normalizing constant $C_{p,g}$, is proportional to

$$\mathcal{L}(\mathbf{\Sigma}_1, \ldots, \mathbf{\Sigma}_k) = \sum_{k=1}^{K} \pi_k \mathcal{L}_k(\mathbf{\Sigma}_k), \tag{4}$$

where

$$\mathcal{L}_k(\mathbf{\Sigma}_k) = \frac{1}{n_k} \sum_{i=1}^{n_k} \rho_k(\mathbf{x}_{ki}^{\mathrm{T}} \mathbf{\Sigma}_k^{-1} \mathbf{x}_{ki}) - \log |\mathbf{\Sigma}_k^{-1}|, \tag{5}$$

and $\rho_k(t) = -2 \log g_k(t)$. The nature of M-estimation is to then divorce the estimators obtained from minimizing (4) from the distributions that generated the negative log-likelihood function. When using the sample covariance matrix, for example, one does need to assume it is based on a sample from a multivariate normal distribution or even from an elliptical distribution. In general, for respective loss functions $\rho_k : \mathbb{R}^+ \to \mathbb{R}^+$, not necessarily related to any g_k, a minimizer (4) represents an M-estimator of scatter. For more detail discussions on the concepts underlying M-estimation and other robust methods, see [3, 17, 18].

Minimizing (4) over $\mathbf{\Sigma}_1, \ldots, \mathbf{\Sigma}_K \in \mathcal{S}(p)$ is equivalent to minimizing (5) individually over $\mathbf{\Sigma}_k \in \mathcal{S}(p)$ for $k = 1, \ldots, K$, i.e. obtaining the individual M-estimators of scatter for each group. One drawback to this approach is that the individual M-estimators of scatter do not exist when $n_k < p$ [19], and do not differ substantially from the sample covariance matrix when n_k is only slightly larger than p. Consequently, for sparse group data, we need to pool the information in the different groups and hence presume that the scatter matrices are somewhat similar across the groups. The most extreme and most common assumption is that the scatter matrices are equal across groups. Here, though, we make no strong model assumptions regarding the different scatter matrices, but rather propose the following two penalization approaches.

Proposal 1: Regularization towards a pooled scatter matrix. A pooled M-estimator of scatter, obtained by pooling together the data from each of the K groups, can be defined as a minimum of

$$\mathcal{L}(\mathbf{\Sigma}) = \sum_{k=1}^{K} \pi_k \mathcal{L}_k(\mathbf{\Sigma}) = \frac{1}{N} \left\{ \sum_{k=1}^{K} \sum_{i=1}^{n_k} \rho_k(\mathbf{x}_{ki}^{\mathrm{T}} \mathbf{\Sigma}^{-1} \mathbf{x}_{ki}) \right\} - \log |\mathbf{\Sigma}^{-1}|. \tag{6}$$

over $\mathbf{\Sigma} \in \mathcal{S}(p)$. Penalized M-estimators of scatter for the individual groups can then be defined as a solution to the optimization problem

$$\min_{\mathbf{\Sigma}_k \in \mathcal{S}(p)} \left\{ \mathcal{L}_k(\mathbf{\Sigma}_k) + \lambda \, d(\mathbf{\Sigma}_k, \hat{\mathbf{\Sigma}}) \right\}, \quad k = 1, \ldots K, \tag{7}$$

where $\hat{\mathbf{\Sigma}}$ is minimizer of (6), $d(\mathbf{\Sigma}_k, \mathbf{\Sigma})$ represents a penalty based on distances between $\mathbf{\Sigma}_k$ and $\mathbf{\Sigma}$, and λ is positive tuning parameter, chosen by the user, which balance the interplay between unrestricted M-estimation of scatter and

shrinkage towards $\mathbf{\Sigma}$. Equivalently, we can write optimization program in (7) in the form

$$\min_{\mathbf{\Sigma}\in\mathcal{S}(p)} \left\{\beta\mathcal{L}_k(\mathbf{\Sigma}_k) + (1-\beta)d(\mathbf{\Sigma}_k, \hat{\mathbf{\Sigma}})\right\}, \tag{8}$$

where penalty parameter $\beta \in (0,1]$ is one-to-one with $\lambda > 0$ via mapping $\lambda = (1-\beta)/\beta$. Formulation (7) is in many ways more instructive as it depicts the role of the penalty term in more lucid manner: one may view the penalty parameter β as a "probability" or degree of belief one assigns on the cost function $\mathcal{L}_k(\mathbf{\Sigma}_k)$ relative to the penalty term $d(\mathbf{\Sigma}_k, \hat{\mathbf{\Sigma}})$. Moreover, β is conveniently on scale $(0,1]$. The latter formulation (8) via regularization parameter β will be used when constucting the fixed point algorithms in Section 5. Examples of penalty functions $d(\mathbf{\Sigma}_k, \mathbf{\Sigma})$ and their properties are addressed in Section 4.

Proposal 2: Joint regularization enforcing similarity among the group scatter matrices. Rather than first defining a pooled scatter matrix, our second proposal simultaneously estimates the group scatter matrices $\mathbf{\Sigma}_k$ along with their 'center' $\mathbf{\Sigma}$. The optimization program is now

$$\underset{\{\mathbf{\Sigma}_k\}_{k=1}^K, \mathbf{\Sigma}\in\mathcal{S}(p)^p}{\text{minimize}} \sum_{k=1}^K \pi_k \left\{\mathcal{L}_k(\mathbf{\Sigma}_k) + \lambda\, d(\mathbf{\Sigma}_k, \mathbf{\Sigma})\right\}. \tag{9}$$

The penalty term $d(\mathbf{\Sigma}_k, \mathbf{\Sigma})$ is as before, but now is viewed as enforcing similarity among the $\mathbf{\Sigma}_k$-s, and the 'center' $\mathbf{\Sigma}$ is now viewed as an 'average' of the $\mathbf{\Sigma}_k$-s. Note again that it is possible to write (9) via penalty parameter $\beta \in (0,1]$ (where $\beta = 1/(1+\lambda)$) as in (8) in which case the term $\mathcal{L}_k(\mathbf{\Sigma}_k) + \lambda\, d(\mathbf{\Sigma}_k, \mathbf{\Sigma})$ in (9) is replaced by $\beta\mathcal{L}_k(\mathbf{\Sigma}_k) + (1-\beta)\, d(\mathbf{\Sigma}_k, \mathbf{\Sigma})$. Note that for fixed $\mathbf{\Sigma}_1, \ldots, \mathbf{\Sigma}_K$, the value of $\mathbf{\Sigma}$ is given by

$$\mathbf{\Sigma}(\boldsymbol{\pi}) = \underset{\mathbf{\Sigma}\in\mathcal{S}(p)}{\arg\min} \sum_{i=1}^K \pi_k\, d(\mathbf{\Sigma}_k, \mathbf{\Sigma}), \tag{10}$$

which represents the *weighted mean* associated with the distance d. For example, the Euclidean, or Frobenius, distance $d_{\mathrm{F}}(\mathbf{\Sigma}_k, \mathbf{\Sigma}) = \left\{\mathrm{Tr}[(\mathbf{\Sigma}_k - \mathbf{\Sigma})^2]\right\}^{1/2}$ gives the standard weighted *arithmetic mean* $\mathbf{\Sigma}_{\mathrm{F}}(\boldsymbol{\pi}) = \sum_{k=1}^K \pi_k \mathbf{\Sigma}_k$.

Modest modifications to Proposals 1 and 2 can be considered. For example, one might consider replacing the tuning constant λ in either proposal with individual tuning constants, say λ_k, $k = 1, \ldots, K$. Typically one tends to choose a larger tuning constant when sample sizes are smaller. However, in our proposals, this does not seem to be necessary since for a particular group, say group j, for which n_j is the smallest, the term $d(\mathbf{\Sigma}_j, \mathbf{\Sigma})$, in either proposal, affects $\mathbf{\Sigma}_j$ more then the other groups since group j affects the value of $\mathbf{\Sigma}$ the least. Another modification to proposal 1 is to consider other pooled estimates of scatter. In particular, if the total sample size N is small, and in particular if $N < p$, then we recommend adding a penalty term to (6) itself, say one which penalized $\mathbf{\Sigma}$ for deviations from $\mathbf{I}$ or deviations from proportionality to $\mathbf{I}$, see e.g., [13] or [14]. We also recommend such an additional penalty term to (9) in Proposal 2 when N is relatively small.

For the special case, $\rho_k(t) = t$ for $k = 1, \ldots K$, the solution for $\mathbf{\Sigma}_k$ in Proposal 1 is

$$\mathbf{S}_k(\beta) = \beta \mathbf{S}_k + (1 - \beta)\mathbf{S}, \tag{11}$$

where $\mathbf{S}_k = \frac{1}{n_k} \sum_{i=1}^{n_k} \mathbf{x}_{ki} \mathbf{x}_{ki}^\top$ is the sample covariance matrix for the kth group, $\mathbf{S} = \sum_{k=1}^{K} \pi_k \mathbf{S}_k$ is the pooled sample covariance matrix, and $\beta = 1/(1+\lambda)$. Note that as the tuning constant $\lambda \to \infty$, $\mathbf{S}_k(\beta) \to \mathbf{S}$, and as $\lambda \to 0$, $\mathbf{S}_k(\beta) \to \mathbf{S}_k$. The estimator (11) is the one proposed by Friedman in [1] in the context of regularized discriminant analysis. Hence, Proposal 1 can be view as a direct generalization of Friedman's estimator.

2.2. Examples of loss functions

Throughout, we assume the loss functions $\rho_k(t)$, $k = 1, \ldots K$, satisfy the following condition:

Condition 1. *The loss functions $\rho_k(t)$, $k = 1, \ldots, K$ are nondecreasing and continuous for $0 < t < \infty$. In addition, $\rho_k(t)$ is convex in $\log t$, i.e. the function $r_k(x) = \rho_k(e^x)$ is convex for $-\infty < x < \infty$.*

Typically, the loss functions $\rho_k(t)$ will be the same for $k = 1, \ldots, K$, but our general development allows for the case when they may differ. Also, the loss functions are often standardized so that the estimators obtained by minimizing (5) are Fisher consistent when the kth sample represents a random sample from the Gaussian distribution $\mathcal{N}_p(\mathbf{0}, \mathbf{\Sigma}_k)$. This holds if and only if $E[\psi_k(\chi^2_p)] = p$, where $\psi_k(t) = t u_k(t)$ and $u_k(t) = \rho_k'(t)$,

Below we provide some common examples of loss functions ρ_k often encountered in the literature and used in multivariate analysis, along with their corresponding weight functions u_k. The weight functions themselves are needed in section 5 to represent the corresponding M-estimating equations and in deriving fixed-point algorithms for the estimators.

(i) **Gaussian loss function**. The density generator for $\mathcal{N}_p(\mathbf{0}, \mathbf{\Sigma})$ is $g(t) = \exp(-t/2)$. Hence, the corresponding loss and weight functions are $\rho_{\mathrm{G}}(t) = t$ and $u_G(t) = 1$ respectively. The corresponding objective function for the kth sample, i.e. (5), is then

$$\mathcal{L}_{\mathrm{G},k}(\mathbf{\Sigma}_k) = \mathrm{Tr}(\mathbf{\Sigma}_k^{-1}\mathbf{S}_k) - \log|\mathbf{\Sigma}_k^{-1}| \tag{12}$$

where $\mathbf{S}_k$ again denotes the sample covariance matrix of the k-th sample.

(ii) t **loss functions**: The density generator for a p-variate elliptical t-distribution on $\nu > 0$ degrees of freedom is $g_\nu(t) = (\nu + t)^{-\frac{1}{2}(\nu+p)}$. Hence, the corresponding loss and weight functions are $\rho_\nu(t) = (\nu + p)\log(\nu + t)$ and $u_\nu(t) = (\nu + p)/(\nu + t)$ respectively. The resulting M-estimators of scatter are not Fisher consistent at a multivariate Gaussian distribution. However, one can obtain such a Fisher consistent version of the t M-estimators by taking the loss function to be $\rho_{\nu,b}(t) \equiv \rho_\nu(t)/b$, with b chosen so that $b = E[\psi_\nu(\chi^2_p)]/p$ and where $\psi_\nu(t) = t u_\nu(t)$. This gives $b = \{(\nu + p)/p\}E[\chi^2_p/(\nu + \chi^2_p)]$.

(iii) **Huber's loss function**: In his seminal work, Huber [20] proposed a family of univariate heavy-tailed distributions often referred to as "least favourable distributions" (LFDs). A LFD corresponds to a symmetric unimodal distribution which follows a Gaussian distribution in the middle, and a double exponential distribution in the tails. The corresponding maximum likelihood estimators are then referred to as Huber's M-estimators. The extension of Huber's M-estimators to the multivariate setting, is usually defined as a generalization of the corresponding univariate M-estimating equations to the multivariate setting, see e.g., [4].

Here, we illustrate how Huber's M-estimators of multivariate scatter can be viewed as maximum likelihood estimators for a family of heavy-tailed p-variate elliptical distributions, namely those with density generator of the form $g_{\mathrm{H}}(t;c) = \exp\{-(1/2)\rho_{\mathrm{H}}(t;c)\}$,where

$$\rho_{\mathrm{H}}(t;c) = \begin{cases} t/b & \text{for } t \leqslant c^2, \\ (c^2/b)\big(\log(t/c^2)+1\big) & \text{for } t > c^2. \end{cases} \tag{13}$$

These distributions follow a multivariate Gaussian distribution in the middle, but have tails that die down at an inverse polynomial rate. The distribution is a valid distribution for $c > 0$, and for the corresponding maximum likelihood estimator of scatter, i.e. the Huber M-estimator of multivariate scatter, the index c represents a user defined tuning constant that determines the robustness and efficiency of the estimator. The constant $b > 0$ represents a scaling factor since it has the effect that if $\widehat{\mathbf{\Sigma}}$ represents the resulting Huber's M-estimator of scatter whenever $b = 1$, the Huber's M estimator of scatter when $b = b_o$ is simply $b_o\widehat{\mathbf{\Sigma}}$. The scaling constant b is usually chosen so that the resulting scatter estimator of scatter is Fisher consistent for the covariance matrix at a chosen reference p-variate elliptical distribution, commonly the p-variate Gaussian distribution. Given a value of c, the value of b needed to obtain Fisher consistency at Gaussian distributions is $b = F_{\chi^2_{p+2}}(c^2) + c^2(1 - F_{\chi^2_p}(c^2))/p$.

We refer to $\rho_{\mathrm{H}}(t;c)$ as Huber's loss function, since it gives rise to Huber's weight function, namely

$$u_{\mathrm{H}}(t;c) = \rho'_{\mathrm{H}}(t;c) = \begin{cases} 1/b, & \text{for } t \leqslant c^2 \\ c^2/(tb), & \text{for } t > c^2 \end{cases}.$$

Thus, an observation $\mathbf{x}$ with squared Mahalanobis distance (MD) $t = \mathbf{x}^\top\mathbf{\Sigma}^{-1}\mathbf{x}$ smaller than c^2 receives constant weight, while observations with large MD are heavily downweighted.

(iv) **Tyler's loss function**: The Gaussian loss function can be viewed as a limiting case of either a t loss function or Huber's loss function by considering $\nu \to \infty$ or $c \to \infty$ respectively. At the other extreme, i.e. as $\nu \to 0$ or $c \to 0$, one obtains Tyler's loss function $\rho_{\mathrm{T}}(t) = p\log t$, whose corresponding weight function is $u_{\mathrm{T}}(t) = p/t$. To obtain this limit using Huber's loss function, first note that the Huber's M-estimator is not affected by replacing $\rho_{\mathrm{H}}(t;c)$ with $\rho^*_{\mathrm{H}}(t;c) = \rho_{\mathrm{H}}(t;c) - h(c,b)$, with $h(c,b) = c^2\{1-\log(c^2)\}/b$ being constant in t.

Then, since $c^2/b \to p$ as $c \to 0$, it follows that $\rho^*_{\mathrm{H}}(t;c) \to \rho_{\mathrm{T}}(t)$. Using this loss function, the corresponding objective function (5) for the kth sample becomes

$$\mathcal{L}_{\mathrm{T},k}(\mathbf{\Sigma}_k) = \frac{p}{n_k}\sum_{i=1}^{n_k}\log\left(\mathbf{x}_{ki}^\top\mathbf{\Sigma}_k^{-1}\mathbf{x}_{ki}\right) - \log|\mathbf{\Sigma}_k^{-1}|. \tag{14}$$

A minimizer of (14) yields Tyler's [5] distribution-free M-estimator of scatter. Note that (14) does not have a unique minimum, since if $\mathbf{\Sigma}_k$ is a minimum then so is $b\mathbf{\Sigma}_k$ for any $b > 0$. That is, Tyler's M-estimator estimates the shape of $\mathbf{\Sigma}_k$ only. A Fisher consistent estimator of the covariance matrix at a Gaussian distribution can be obtained by multiplying any particular minimum $\mathbf{\Sigma}_k$ by $b_k = \mathrm{Median}\{\mathbf{x}_{ki}^\top\mathbf{\Sigma}_k^{-1}\mathbf{x}_{ki}; i = 1,\ldots,n_k\}/\mathrm{Median}(\chi^2_p)$. In discriminant application reported in Section 7, this scaling is utilized.

It is worth noting that the objective function (14) does not correspond to the negative log-likelihood of any family of RES distributions, since $g(t) = e^{-\rho_{\mathrm{T}}(t)/2} = t^{-p/2}$ is not a valid density generator. However, (14) does correspond the negative log-likelihood function for a p-variate Angular Central Gaussian (ACG) distribution [21]. The ACG distribution is defined on the unit p-sphere $\mathcal{S}^{p-1} = \{\theta \in \mathbb{R}^p; \theta^\top\theta = 1\}$, and its p.d.f. relative to the uniform distribution on $\mathcal{S}^{p-1}$ has the form

$$f(\mathbf{x}|\mathbf{\Sigma}) = C_p|\mathbf{\Sigma}|^{-1/2}(\mathbf{x}^\top\mathbf{\Sigma}^{-1}\mathbf{x})^{-p/2}, \quad \mathbf{x} \in \mathcal{S}^{p-1}. \tag{15}$$

Here the scatter matrix parameter $\mathbf{\Sigma} \in \mathcal{S}(p)$ is uniquely defined up to a positive scalar. Although ACG distribution does not belong to the class of RES distributions, it is related to it. Namely, an important property of the elliptical family is that $\mathbf{x}/\|\mathbf{x}\|$ has ACG distribution for any $\mathbf{x} \sim \mathcal{E}_p(0,\mathbf{\Sigma},g)$. Note that replacing $\mathbf{x}_{ki}$ with $\mathbf{x}_{ki}/\|\mathbf{x}_{ki}\|$ in (14) does not affect its minimizer since it is equivalent to subtracting the term $\frac{p}{n_k}\sum_{i=1}^{n_k}\log\left(\mathbf{x}_{ki}^\top\mathbf{x}_{ki}\right)$, which does not depend on $\mathbf{\Sigma}_k$. Consequently, the distribution of the resulting M-estimator of scatter is the same under any elliptical distribution.

3. Preliminaries on g-convexity

The optimizaton problems defined by (6), (7) and (9) are easiest to handle when the target functions to be minimized are convex. The Gaussian negative log-likelihood (12), for example, is well known to be strictly convex as a function of the inverse covariance matrix. Unfortunately, the functions (6), (7) and (9) in general tend not to possess this convexity property. Other notions of convexity, though, can be applied. Briefly summarizing, convexity properties of sets in metric spaces depend on the definition of the shortest paths (geodesic curves) between pairs of points. Thus, when the metric is altered, geodesic curves change and consequently so does the notion of convexity. In our treatment, we use the notion of g-convexity relative to the intrinsic Riemannian manifold structure of the positive semi-definite cone; see [22] and [14] for a more detailed exposition.

The set $\mathcal{S}(p)$ can be endowed with a smooth Riemannian manifold structure by changing the usual Euclidean metric to the Riemannian one. The latter

can be defined by stipulating the notion of a *geodesic path* from $\mathbf{\Sigma}_0 \in \mathcal{S}(p)$ to $\mathbf{\Sigma}_1 \in \mathcal{S}(p)$ and setting it to be

$$\mathbf{\Sigma}_t = \mathbf{\Sigma}_0^{1/2} \left(\mathbf{\Sigma}_0^{-1/2} \mathbf{\Sigma}_1 \mathbf{\Sigma}_0^{-1/2} \right)^t \mathbf{\Sigma}_0^{1/2}, \quad t \in [0,1]. \tag{16}$$

Given $\mathbf{\Sigma}_0, \mathbf{\Sigma}_1 \in \mathcal{S}(p)$, we have $\mathbf{\Sigma}_t \in \mathcal{S}(p)$ for $0 \leqslant t \leqslant 1$, therefore, $\mathcal{S}(p)$ is a *geodesically convex set*. A function $h : \mathcal{S}(p) \to \mathbb{R}$ is a *g-convex function* if

$$h(\mathbf{\Sigma}_t) \leqslant (1-t)\ h(\mathbf{\Sigma}_0) + t\ h(\mathbf{\Sigma}_1), \ \ t \in (0,1). \tag{17}$$

If the inequality is strict, h is said to be strictly g-convex. When $p = 1$, g-convexity/strict g-convexity is equivalent to the function $h(e^x)$ being convex/strictly convex over $\mathbb{R}$, i.e. $h(s)$ is convex/strictly convex in $\log(s)$ for $s > 0$. the concept of g-convexity enjoys properties similar to those of convexity in Euclidean spaces. In particular, if h is g-convex on $\mathcal{S}(p)$ then any local minimum is a global minimum. Furthermore, if a minimum is obtained in $\mathcal{S}(p)$ then the set of all minima form a g-convex subset of $\mathcal{S}(p)$. If h is strictly g-convex and a minimum is obtained in $\mathcal{S}(p)$, then it is unique, see [9, 14] and reference therein for more details. An important additional property of g-convexity, not shared by convexity in Euclidean spaces, is that if $h(\mathbf{\Sigma})$ is g-convex/strict g-convex in $\mathbf{\Sigma}$, then it is also g-convex/strict g-convex in $\mathbf{\Sigma}^{-1}$.

Similarly, given a naturally induced manifold structure over $\mathcal{S}(p) \times \mathcal{S}(p)$, we say that a function $u : \mathcal{S}(p) \times \mathcal{S}(p) \to \mathbb{R}$ is jointly g-convex if

$$u(\mathbf{\Sigma}_t^*, \mathbf{\Sigma}_t) \leqslant (1-t)\ u(\mathbf{\Sigma}_0^*, \mathbf{\Sigma}_0) + t\ u(\mathbf{\Sigma}_1^*, \mathbf{\Sigma}_1) \ \text{ for } \ t \in (0,1),$$

where $\mathbf{\Sigma}_t^*$ and $\mathbf{\Sigma}_t$ are defined as in (16).

To establish the existence and uniqueness of the solutions to the minimization problems (6), (7) and (9), a basic requirement is that loss functions $\mathcal{L}_k(\mathbf{\Sigma}_k)$ in (5) be g-convex in $\mathbf{\Sigma}_k \in \mathcal{S}(p)$. This is achieved when the respective loss functions $\rho_k(t)$, $k = 1, \ldots, K$, satisfy Condition 1, see [11] or [13, Lemma 1]. Many common loss functions satisfy this condition, such as the Gaussian, t, Huber and Tyler loss functions given in Section 2.2. Using the terminology of this section, Condition 1 simply requires that $\rho_k(t)$ be a nondecreasing, continuous g-convex function.

4. Distance measures for covariance matrices

Optimization problems (7) and (9) balance the overall loss between the separate group M-estimation with shrinkage towards a mutual joint (pooling) center. The penalty terms introduced in (7) and (9) require specifying a distance function $d(\mathbf{A}, \mathbf{B}) : \mathcal{S}(p) \times \mathcal{S}(p) \to \mathbb{R}_0^+$, (we sometimes refer to d as a penalty function due to its role here). Listed below are some properties one may desire for a distance function.

(D1) $d(\mathbf{A}, \mathbf{B}) = 0$ if $\mathbf{A} = \mathbf{B}$,

(D2) $d(\mathbf{A}, \mathbf{B})$ is jointly g-convex,

(D3) *symmetry*: $d(\mathbf{A}, \mathbf{B}) = d(\mathbf{B}, \mathbf{A})$.

(D4) *affine invariance*: $d(\mathbf{A}, \mathbf{B}) = d(\mathbf{C}\mathbf{A}\mathbf{C}^\top, \mathbf{C}\mathbf{B}\mathbf{C}^\top)$ for any nonsingular $\mathbf{C}$.

(D5) *scale invariance*: $d(c_1\mathbf{A}, c_2\mathbf{B}) = d(\mathbf{A}, \mathbf{B})$ for $c_1, c_2 > 0$,

D1 and D2 are necessary requirements. D1 is an obvious requirement, whereas D2 is needed to guarantee that the optimization problem (9) is g-convex. Properties D3, D4 and D5 are considered optional. When property D4 holds, the resulting estimators of the scatter matrices are affine equivariant. That is, if we transform the data $\mathbf{x}_{ki} \to \mathbf{C}\mathbf{x}_{ki}$ for all $k = 1, \ldots, K$; $i = 1, \ldots, n_k$, then $\{\boldsymbol{\Sigma}_1, \ldots, \boldsymbol{\Sigma}_K, \boldsymbol{\Sigma}\} \to \{\mathbf{C}\boldsymbol{\Sigma}_1\mathbf{C}^\top, \ldots, \mathbf{C}\boldsymbol{\Sigma}_K\mathbf{C}^\top, \mathbf{C}\boldsymbol{\Sigma}\mathbf{C}^\top\}$. Property D5 is useful if we are primarily interested in the shape of the scatter matrices, that is, the scatter matrices up to a positive scalar. Some scatter scatter estimators (such as Tyler's M-estimator) are shape estimators only, and for such shape estimators, D5 is necessary. Property D5 is also important if the individual covariances are believed to be scaled differently, and so one may wish to only poll together their shapes but not their overall scale.

As noted in section 2.1, every distance function induces a notion of a mean, defined by (10). For example, when $\mathcal{S}(p)$ is treated as being embedded within the space of symmetric matrices, it inherits the natural Euclidean distance $d_{\mathrm{F}}(\mathbf{A}, \mathbf{B}) = \|\mathbf{A} - \mathbf{B}\|_{\mathrm{F}}^2$, which is the usual Frobenius metric. The mean (10) corresponding to the Euclidean distance d_{F} is the weighted *arithmetic mean*

$$\boldsymbol{\Sigma}_{\mathrm{F}}(\boldsymbol{\pi}) = \sum_{k=1}^{K} \pi_k \boldsymbol{\Sigma}_k. \tag{18}$$

Unfortunately, the Frobenius metric fails to be jointly g-convex function, and so does not fit into our framework. However, following are examples of g-convex distance functions, which can be used to construct g-convex optimization problems.

4.1. Riemannian distance

The Riemannian distance,

$$d_{\mathrm{R}}(\mathbf{A}, \mathbf{B}) = \|\log(\mathbf{A}^{-1/2}\mathbf{B}\mathbf{A}^{-1/2})\|_{\mathrm{F}}^2,$$

is the length of the geodesic curve between $\mathbf{A}$ and $\mathbf{B}$ as defined in (16) and hence it is a natural and widely studied distance between positive definite matrices. See e.g., [23] and references therein for a comprehensive survey. The Riemannian distance satisfies properties D1, D2, D3 and D4. The unique extremum $\boldsymbol{\Sigma}_{\mathrm{R}}(\boldsymbol{\pi})$ of (10) is a weighted form of the *Karcher mean* (or Riemannian or geodesic mean), see [24, 23]. It was shown in [24] to be the unique positive definite solution of

$$\sum_{k=1}^{K} \pi_k \log(\boldsymbol{\Sigma}_k^{-1}\boldsymbol{\Sigma}_{\mathrm{R}}) = \mathbf{0}. \tag{19}$$

Equation (19) can be written in different forms, e.g., $\sum_{k=1}^{K} \pi_k \log(\mathbf{\Sigma}_{\mathrm{R}}^{1/2}\mathbf{\Sigma}_k^{-1}\mathbf{\Sigma}_{\mathrm{R}}^{1/2}) = \mathbf{0}$, using the formula $\mathbf{A}^{-1}\log(\mathbf{B})\mathbf{A} = \log(\mathbf{A}^{-1}\mathbf{B}\mathbf{A})$, valid for any invertible matrix $\mathbf{A}$ and any matrix $\mathbf{B}$ having real positive eigenvalues [25]. Equation (19) does not have a closed-form solution and a number of rather complex numerical approaches have been proposed to compute the solution. See e.g., [26, 27] and references therein for a number of such techniques.

4.2. Ellipticity distance

The ellipticity distance,

$$d_{\mathrm{E}}(\mathbf{A}, \mathbf{B}) = p \log \frac{1}{p}\mathrm{Tr}(\mathbf{A}^{-1}\mathbf{B}) - \log|\mathbf{A}^{-1}\mathbf{B}|,$$

was first introduced for the penalized robust covariance estimation problem in [10]. Therein, $\mathbf{B}$ corresponds to a fixed shrinkage target shape matrix, and it is shown that $d_{\mathrm{E}}(\mathbf{A}, \mathbf{B})$ is g-convex in $\mathbf{A}$ for fixed $\mathbf{B}$. We refer d_{E} as the *ellipticity distance* since it is related to the ellipticity factor, $e(\mathbf{\Sigma}) = \frac{1}{p}\mathrm{Tr}(\mathbf{\Sigma})/|\mathbf{\Sigma}|^{1/p}$, i.e. the ratio of the arithmetic and geometric means of the eigenvalues of $\mathbf{\Sigma}$. The factor $e(\mathbf{\Sigma}) \geqslant 1$ with equality if and only if $\mathbf{\Sigma} \propto \mathbf{I}$. Its relationship to d_{E} is given by $d_{\mathrm{E}}(\mathbf{A}, \mathbf{B}) = \log e(\mathbf{A}^{-1/2}\mathbf{B}\mathbf{A}^{-1/2})$.

The ellipticity distance is scale invariant, i.e. it satisfies D5. It also satisfies properties D1, D2, and D4. We summarize these properties in the following proposition, which also characterize its induced mean (10). The proof follows readily from the joint g-convexity of the trace term $\mathrm{Tr}(\mathbf{A}^{-1}\mathbf{B})$, which is proven in Lemma 1 of the Appendix.

Proposition 1. The ellipticity distance $d_{\mathrm{E}}(\mathbf{A}, \mathbf{B})$ satisfies D1, D2, D4 and D5. Furthermore, the optimization problem (10) has a unique minimizer (up to a scale), with the minimizer being the unique solution (up to scale) of the fixed-point equation

$$\mathbf{\Sigma}_{\mathrm{E}} = \left(\sum_{k=1}^{K} \pi_k \frac{p\mathbf{\Sigma}_k^{-1}}{\mathrm{Tr}(\mathbf{\Sigma}_k^{-1}\mathbf{\Sigma}_{\mathrm{E}})} \right)^{-1}. \tag{20}$$

The ellipticity induced mean $\mathbf{\Sigma}_{\mathrm{E}}$ is related to the *harmonic mean*, [23], of positive definite matrices. In particular, it can be viewed as an (implicitly) weighted harmonic mean of normalized scatter matrices.

4.3. Kullback-Leilber distance

The information theoretic (Gaussian) Kullback-Leibler (KL) divergence [28] is defined as

$$d_{\mathrm{KL}}(\mathbf{A}, \mathbf{B}) = \mathrm{Tr}(\mathbf{A}^{-1}\mathbf{B}) - \log|\mathbf{A}^{-1}\mathbf{B}| - p.$$

In statistics literature it has gained popularity due to the seminal works of James and Stein [29] who utilized it in the risk function of covariance matrices. It has also been recently used as shrinkage penalty in covariance estimation problems in [30], who considered a single sample case with $\mathbf{B}$ played being a fixed shrinkage target matrix. The next claim shows the g-convexity of KL-distance and provides the respective mean (10).

Proposition 2. The KL-distance $d_{\mathrm{KL}}(\mathbf{A}, \mathbf{B})$ satisfies D1, D2 and D4. Furthermore, the optimization problem (10) has the unique minimizer

$$\boldsymbol{\Sigma}_{\mathrm{KL}}(\boldsymbol{\pi}) = \left(\sum_{k=1}^{K} \pi_k \boldsymbol{\Sigma}_k^{-1} \right)^{-1}, \tag{21}$$

which corresponds to a weighted harmonic mean.

The metrics d_{E} and d_{KL} are thus closely related, however, the former is a scale invariant metric while the latter is not. At the same time $\boldsymbol{\Sigma}_{\mathrm{E}}$ is given by an implicit weighted harmonic mean equation (20) while $\boldsymbol{\Sigma}_{\mathrm{KL}}$ is given explicitly by (21). Thus due to scale invariance, one looses in the simplicity of calculation. The scale invariance property of the penalty function is especially useful in problems when the unknown elliptical distributions of the samples are different (heterogeneous environment) and hence the scatter matrices obtained using same loss function would have a priori different scalings. This follows from the fact that any M-estimator $\hat{\boldsymbol{\Sigma}}$ provides an estimate up to a constant and the constant of proportionality depends on the underlying distribution, as well.

5. Fixed-point algorithms

In this section we propose fixed-point (FP) algorithms for computing the regularized scatter estimators. We first consider Proposal 2, which corresponds to optimization problem (9). If we differentiate (9) with respect to $\boldsymbol{\Sigma}_k^{-1}, k = 1, \ldots, K$ and $\boldsymbol{\Sigma}$, we obtain the first order optimality conditions on the extremum, i.e. the M-estimating equations, which are

$$\frac{\partial \mathcal{L}_k(\boldsymbol{\Sigma}_k)}{\partial \boldsymbol{\Sigma}_k^{-1}} + \lambda \frac{\partial d(\boldsymbol{\Sigma}_k, \boldsymbol{\Sigma})}{\partial \boldsymbol{\Sigma}_k^{-1}} = \mathbf{0}, \quad k = 1, \ldots, K, \tag{22}$$

$$\sum_{k=1}^{K} \pi_k \frac{\partial d(\boldsymbol{\Sigma}_k, \boldsymbol{\Sigma})}{\partial \boldsymbol{\Sigma}} = \mathbf{0}, . \tag{23}$$

If we denote $\partial d(\boldsymbol{\Sigma}_k, \boldsymbol{\Sigma}) / \partial \boldsymbol{\Sigma}_k^{-1}$ by $d'(\boldsymbol{\Sigma}_k, \boldsymbol{\Sigma})$ and use the form of $\mathcal{L}_k(\boldsymbol{\Sigma}_k)$ given in (5), then the first equation becomes

$$\boldsymbol{\Sigma}_k = \frac{1}{n_k} \sum_{i=1}^{n_k} u_k(\mathbf{x}_{ki}^{\mathrm{T}} \boldsymbol{\Sigma}_k^{-1} \mathbf{x}_{ki}) \mathbf{x}_{ki} \mathbf{x}_{ki}^{\mathrm{T}} + \lambda d'(\boldsymbol{\Sigma}_k, \boldsymbol{\Sigma}), \tag{24}$$

with $u_k(t) = \rho_k'(t)$, $k = 1, \ldots, K$, acting as weight functions. The second equation (23) coincides with the definition of the mean given in (10). Hence, a general FP algorithm for finding the solution to the optimization problem (9) is given by the iterative scheme

$$\boldsymbol{\Sigma}_k \leftarrow \frac{1}{n_k} \sum_{i=1}^{n_k} u_k(\mathbf{x}_{ki}^{\mathrm{T}} \boldsymbol{\Sigma}_k^{-1} \mathbf{x}_{ki}) \mathbf{x}_{ki} \mathbf{x}_{ki}^{\mathrm{T}} + \lambda d'(\boldsymbol{\Sigma}_k, \boldsymbol{\Sigma}), \quad k = 1, \ldots, K \tag{25}$$

$$\boldsymbol{\Sigma} \leftarrow \boldsymbol{\Sigma}_k(\boldsymbol{\pi}), \tag{26}$$

which updates the covariance matrices in natural order, i.e., from $\mathbf{\Sigma}_1, \ldots \mathbf{\Sigma}_K$ to $\mathbf{\Sigma}$, and cyclically repeating the procedure until convergence. One may view it as blockwise alternating minimization algorithm in which one minimizes the objective function in one block at a time while keeping others fixed at their current iterates. The convergence properties of such a scheme is omitted and is a subject of a follow up paper (under preparation) in which blockwise minimization majorization (MM) algorithmic scheme [31, 32] is utilized for proving convergence. We note that an MM algorithm (in the single covariance estimation problem, $K = 1$) have been recently used in [30] and [14] in constructing simple convergence proof of FP algorithm of regularized Tyler's M-estimator. At this point it simply suffices to say that in practice, when the scheme converges, then, by g-convexity (or strict g-convexity), it must converge to a solution (or the unique solution) to the optimization problem (9).

For the first proposal, which corresponds to the optimization problem (6)-(7), one first needs to compute $\mathbf{\Sigma}$, i.e. the optimizer of (6). This simply involves computing a non-penalized M-estimator of scatter, for which computational algorithms have been well studied, see e.g., [19, 33]. A simple FP algorithm for $\mathbf{\Sigma}$ is given by the iterative scheme

$$\mathbf{\Sigma} \leftarrow \frac{1}{N} \sum_{k=1}^{K} \sum_{i=1}^{n_k} u_k(\mathbf{x}_{ki}^{\mathrm{T}} \mathbf{\Sigma}^{-1} \mathbf{x}_{ki}) \mathbf{x}_{ki} \mathbf{x}_{ki}^{\mathrm{T}}.$$

Given this value of $\mathbf{\Sigma}$, the FP algorithm for the $\mathbf{\Sigma}_k$'s, the optimizers of (7), corresponds to (25) with $\mathbf{\Sigma}$ held fixed.

The exact forms of the derivatives for $d_{\mathrm{R}}, d_{\mathrm{E}}$ and d_{KL} are respectively:

- $d'_{\mathrm{R}}(\mathbf{\Sigma}_k, \mathbf{\Sigma}) = 2 \log(\mathbf{\Sigma} \mathbf{\Sigma}_k^{-1}) \mathbf{\Sigma}_k,$
- $d'_{\mathrm{E}}(\mathbf{\Sigma}_k, \mathbf{\Sigma}) = \dfrac{p \mathbf{\Sigma}}{\mathrm{Tr}(\mathbf{\Sigma}_k^{-1} \mathbf{\Sigma})} - \mathbf{\Sigma}_k,$
- $d'_{\mathrm{KL}}(\mathbf{\Sigma}_k, \mathbf{\Sigma}) = \mathbf{\Sigma} - \mathbf{\Sigma}_k.$

The specific form of the fixed point (FP) algorithm which utilizing each of these distances is given below. For simplicity, let

$$\mathbf{\Psi}_k(\mathbf{\Sigma}_k) = \frac{1}{n_k} \sum_{i=1}^{n_k} u_k(\mathbf{x}_{ki}^{\mathrm{T}} \mathbf{\Sigma}_k^{-1} \mathbf{x}_{ki}) \mathbf{x}_{ki} \mathbf{x}_{ki}^{\mathrm{T}}.$$

Also, map $\lambda \geq 0$ to $\beta = 1/(1 + \lambda) \in (0, 1]$. Note that β is a regularization parameter given in the formulation (8).

- **Ellipticity distance** $d_{\mathrm{E}}(\mathbf{A}, \mathbf{B})$:

$$\mathbf{\Sigma}_k \leftarrow \beta \mathbf{\Psi}_k(\mathbf{\Sigma}_k) + (1 - \beta) \frac{p \mathbf{\Sigma}}{\mathrm{Tr}(\mathbf{\Sigma}_k^{-1} \mathbf{\Sigma})}, \tag{27}$$

$$\mathbf{\Sigma} \leftarrow \left(\sum_{k=1}^{K} \pi_k \frac{p \mathbf{\Sigma}_k^{-1}}{\mathrm{Tr}(\mathbf{\Sigma}_k^{-1} \mathbf{\Sigma})} \right)^{-1}. \tag{28}$$

Note that in this case the first equation is not sensitive to the scaling of $\mathbf{\Sigma}$, as we expected from the scaling invariance properties of the penalty d_{E}.

- **KL-distance** $d_{\mathrm{KL}}(\mathbf{A}, \mathbf{B})$:

$$\mathbf{\Sigma}_k \leftarrow \beta \mathbf{\Psi}_k(\mathbf{\Sigma}_k) + (1-\beta)\mathbf{\Sigma}, \tag{29}$$

$$\mathbf{\Sigma} \leftarrow \left(\sum_{k=1}^{K} \pi_k \mathbf{\Sigma}_k^{-1} \right)^{-1}. \tag{30}$$

It is shown by Theorems 1 and 2 in [13] that if the loss function ρ_k is bounded below, then for any fixed $\mathbf{\Sigma}$ the FP algorithm (29) always converges to a unique solution.

- **Riemannian distance** $d_{\mathrm{R}}(\mathbf{A}, \mathbf{B})$:

$$\mathbf{\Sigma}_k \leftarrow \left[\mathbf{I} - 2\lambda \log(\mathbf{\Sigma}\mathbf{\Sigma}_k^{-1}) \right]^{-1} \mathbf{\Psi}_k(\mathbf{\Sigma}_k), \tag{31}$$

$$\mathbf{\Sigma} \leftarrow \mathbf{\Sigma}_{\mathrm{R}}(\boldsymbol{\pi}), \tag{32}$$

where $\mathbf{\Sigma}_{\mathrm{R}}(\boldsymbol{\pi})$ is the solution to (19)

Remark 1. Note that iterative algorithms for KL and ellipticity metrics provide simple FP algorithms, but the Riemannian metric does not admit a simple FP equation for the joint center $\mathbf{\Sigma}$ update, but rather requires more complex schemes; see e.g. [27] for an appropriate iterative algorithm. In addition, note that the last step requires solving $\mathbf{\Sigma}_{\mathrm{R}}(\boldsymbol{\pi})$ as a solution to (19) which is computationally demanding task. Also the updates for $\mathbf{\Sigma}_k$, $k = 1, \ldots, K$ in (31) are computationally more demanding than the updates (29) and (27) corresponding to the other penalties. Therefore, since Riemannian distance requires a more specialized algorithm, we do not consider the Riemannian penalty further and exclude it from our simulation studies.

5.1. Tyler's loss function and the ellipticity penalty

In this section we treat Tyler's loss function $\rho_k(t) = p \log t$, for $k = 1, \ldots, K$, in more detail. For this case the M-estimation loss function (5) is scale invariant, i.e., $\mathcal{L}_k(c\mathbf{\Sigma}_k) = \mathcal{L}_k(\mathbf{\Sigma}_k)$, Hence we will mainly consider the scale invariant distance $d_{\mathrm{E}}(\mathbf{A}, \mathbf{B})$ when using Tyler's loss function. Here, the fixed-point iteration in (27) becomes

$$\mathbf{\Sigma}_k \leftarrow \beta \frac{p}{n_k} \sum_{k=1}^{n} \frac{\mathbf{x}_{ik}\mathbf{x}_{ik}^{\mathrm{T}}}{\mathbf{x}_{ik}^{\mathrm{T}}\mathbf{\Sigma}_k^{-1}\mathbf{x}_{ik}} + (1-\beta)\frac{p}{\mathrm{Tr}(\mathbf{\Sigma}_k^{-1}\mathbf{\Sigma})}\mathbf{\Sigma}, \tag{33}$$

with the joint center update being the same as in (28). The resulting estimators $\mathbf{\Sigma}_1, \ldots, \mathbf{\Sigma}_K$ and $\mathbf{\Sigma}$ are well defined only up to their shape. That is, $\{\mathbf{\Sigma}_1, \ldots, \mathbf{\Sigma}_K, \mathbf{\Sigma}\}$ is a solution if and only if $\{\sigma_1\mathbf{\Sigma}_1, \ldots, \sigma_K\mathbf{\Sigma}_K, \sigma\mathbf{\Sigma}\}$ is a solution for any positive $\sigma_1, \ldots, \sigma_K$ and σ.

Curiously, if we choose a solution to (33) and (28) for which $\mathrm{Tr}(\mathbf{\Sigma}_k^{-1}\mathbf{\Sigma}) = p$, then we see that this also gives the solution to (29) and (30), i.e. when using Tyler's loss function with KL-penalization. As noted previously, for fixed $\mathbf{\Sigma}$, a unique solution to (29) always exists whenever the loss function ρ_k is bounded below. Tyler's loss function, though, is not bounded below, and additional condition are needed to ensure existence. In particular, as a corollary to Theorem 4 in [13] we have the following result.

Theorem 1. For Tyler's loss function $\rho_k(t) = p\log t$ and $0 \leqslant \beta = 1/(1+\lambda) < 1$, a necessary condition for program (9) to have a non-singular minimum is that for each of the $k = 1, \ldots, K$ group samples the inequality

$$P_{n_k,k}(\mathcal{V}) = \frac{\#\{\mathbf{x}_{ik} \in \mathcal{V}\}}{n_k} < \frac{\dim(\mathcal{V})}{p\beta} \tag{34}$$

holds for any subspace $\mathcal{V}$ of $\mathbb{R}^p$. Furthermore, if we replace the '$<$' with '$\leq$' in Condition (34), then this becomes a sufficient condition for ensuring, for fixed $\mathbf{\Sigma}$, that (33) admit unique solutions for $\mathbf{\Sigma}_k$ up to a scale.

Here, $P_{n_k,k}$ is the empirical measure for the k-th group sample. Condition (34) implies that when the data is in general position for each group sample, we need $n_k > p\beta$.

6. Cross validation procedure

Let us describe a simple cross validation (CV) procedure that can be utilized for penalty parameter selection $\beta \in (0,1)$. Recall that the objective function given all the data for the parameters $\mathbf{\Sigma}_1, \ldots, \mathbf{\Sigma}_K$ is

$$\mathcal{L}(\mathbf{\Sigma}_1, \ldots, \mathbf{\Sigma}_K) = \sum_{k=1}^{K} \left\{ \sum_{i=1}^{n_k} \rho_k(\mathbf{x}_{ki}^\top \mathbf{\Sigma}_k^{-1} \mathbf{x}_{ki}) - n_k \log|\mathbf{\Sigma}_k^{-1}| \right\}.$$

Partition each data set $\mathbf{X}_k = \{\mathbf{x}_{k1}, \ldots, \mathbf{x}_{kn_k}\}$ into Q separate sets of approximately similar size (or exactly equal size when $\mathrm{mod}(n_k, Q) = 0$), i.e., let $I_{k1} \cup I_{k2} \cup \cdots \cup I_{kQ} = \{1, \ldots, n_k\} \equiv [n_k]$ denote the indices of Q data folds of the kth data set. Common choises are $Q = 5, 10$ or $Q = n_k$, which is known as *leave-one-out cross validation*. When we leave qth fold out from the kth data set $\mathbf{X}_k$, we obtain a reduced data set, denoted by data set $\mathbf{X}_{-q,k}$, that does not include the observations $\{\mathbf{x}_{kq}\}$, $q \in [n_k] \setminus I_{kq}$, in the qth fold. Cross validation scheme then proceeds as follows:

1. **for** $\beta \in [\beta]$ (= a grid of β values in $(0,1)$) and $q \in \{1, \ldots, Q\}$ **do**
   - Compute $\big\{\hat{\mathbf{\Sigma}}_k(\beta, q)\big\}_{k=1}^K$ based on the data sets $\big\{\mathbf{X}_{-q,k}\big\}_{k=1}^K$.
   - *CV fit* for β is computed over the qth folds that were left out:

$$\mathrm{CV}(\beta, q) = \sum_{k=1}^{K} \left\{ \sum_{\tilde{q} \in I_{kq}} \rho_k\big(\mathbf{x}_{k\tilde{q}}^\top \big[\hat{\mathbf{\Sigma}}_k(\beta, q)\big]^{-1} \mathbf{x}_{k\tilde{q}}\big) - (\#I_{kq}) \cdot \log\big|\hat{\mathbf{\Sigma}}_k(\beta, q)^{-1}\big| \right\} \tag{35}$$

   where $\#I_{kq}$ denotes the cardinality of set I_{kq}.

end

2. Compute the average CV fit: $\mathrm{CV}(\beta) = (1/Q)\sum_{q=1}^{Q}\mathrm{CV}(\beta, q)$, $\forall\beta \in [\beta]$.

3. Select $\hat{\beta}_{\mathrm{CV}} = \arg\min_{\beta\in[\beta]}\mathrm{CV}(\beta)$.

4. Compute $\left\{\hat{\mathbf{\Sigma}}_k(\beta)\right\}_{k=1}^{K}$ based on the entire data sets $\{\mathbf{X}_k\}_{k=1}^{K}$ for $\beta = \hat{\beta}_{\mathrm{CV}}$.

It is easy to imagine a variant of this approach in which definition of CV fit is tuned towards a measure that arises from application perspective. For example, in discriminant analysis described in Section 7 one may wish to replace the CV fit measure in (35) by the classification error rate over the qth folds. This approach, however, gave essentially same results, and hence the CV scheme described above is used also in this setting due to its simplicity.

7. Regularized discriminant analysis (RDA)

The classic Fisher's QDA is based on the assumption that each class contain a sample of i.i.d. random vectors from the p-variate Gaussian distribution with mean vector $\boldsymbol{\mu}_k$ and covariance matrix $\mathbf{\Sigma}_k$. For simplicity of exposition we assume that the class prior probabilities are equal. The QDA classification rule then assigns a new measurement $\mathbf{x}$ to a group $\hat{k}$, where

$$\hat{k} = \arg\min_{1\le k\le K}\left\{(\mathbf{x}-\boldsymbol{\mu}_k)^\top\mathbf{\Sigma}_k^{-1}(\mathbf{x}-\boldsymbol{\mu}_k) + \ln|\mathbf{\Sigma}_k|\right\}. \tag{36}$$

If all class covariance matrices are presumed to be identical, i.e., $\mathbf{\Sigma}_k = \mathbf{\Sigma}$ for $k = 1,\ldots,K$, then the rule simplifies to $\hat{k} = \arg\min_{1\le k\le K}(\mathbf{x}-\boldsymbol{\mu}_k)^\top\mathbf{\Sigma}^{-1}(\mathbf{x}-\boldsymbol{\mu}_k)$, referred to as LDA rule hereafter. In general, QDA or LDA perform well when the class distributions are approximately normal and good estimates based on the training data can be obtained for the population parameters, mean vectors $\boldsymbol{\mu}_k$ and covariance matrices $\mathbf{\Sigma}_k$. These are usually estimated by the sample mean vectors, $\bar{\mathbf{x}}_k = (1/n_k)\sum_{i=1}^{n_k}\mathbf{x}_{ki}$, and sample covariance matrices $\mathbf{S}_k = \frac{1}{n_k}\sum_{i=1}^{n_k}(\mathbf{x}_{ki}-\bar{\mathbf{x}}_k)(\mathbf{x}_{ki}-\bar{\mathbf{x}}_k)^\top$ of the training samples $\mathbf{X}_k$, $k = 1,\ldots,K$. QDA generally requires larger sample sizes than LDA and is often reported to be more sensitive to violations of the assumptions. QDA also can not be applied if $n_k \le p$ for any class and may exhibit poor performance when n_k is not considerably larger than the dimension p. LDA has the benefit of requiring only that $N = \sum_{k=1}^{K} n_k > p$. LDA can be viewed as a form of regularized QDA that decreases the variance by using a pooled covariance matrix estimate, $\mathbf{S} = \sum_{k=1}^{K}\pi_k\mathbf{S}_k$. This can sometimes lead to superior performance compared to QDA especially in small-sample settings even if the population class covariance matrices are substantially different.

The idea in RDA proposed in [1] is to replace the unknown covariance matrices $\mathbf{\Sigma}_k$ in the QDA rule (36) by shrinkage estimates $\mathbf{S}_k(\beta)$ defined in (11), where $\beta \in [0,1]$ denotes the shrinkage regularization parameter. If $\beta = 1$, then one obtains the conventional empirical QDA rule and if $\beta = 0$, then one obtains

the empirical LDA rule based on the pooled sample covariance matrix $\mathbf{S}$. A value $\beta \in (0,1)$, between these two extremes, then offers a compromise between LDA and QDA. In our RDA approach we use the developed robust estimators $\hat{\mathbf{\Sigma}}_k(\beta)$ instead of the shrinkage sample covariance matrices $\mathbf{S}_k(\beta)$. The RDA rule becomes

$$\hat{k} = \arg \min_{1 \le k \le K} \left\{ (\mathbf{x} - \hat{\boldsymbol{\mu}}_k)^\top [\hat{\mathbf{\Sigma}}_k(\beta)]^{-1} (\mathbf{x} - \hat{\boldsymbol{\mu}}_k) + \ln |\hat{\mathbf{\Sigma}}_k(\beta)| \right\}. \tag{37}$$

For robust loss functions, we employ the spatial median [34] as an estimate $\hat{\boldsymbol{\mu}}_k$ of location, whereas sample mean is used for Gaussian loss function. Note that the shrinkage scatter matrix estimators $\hat{\mathbf{\Sigma}}(\beta)$ are computed using the centered data.

7.1. Simulation set-up

Population class conditional distributions are chosen to be p-variate elliptical distributions and the total sample size is fixed to $N = \sum_{k=1}^{K} n_k = 100$, the number of groups is $K = 3, 5$ and dimension varies from $p = 10, 20, 30$. For simplicity we use the same loss function $\rho = \rho_k$ for each K samples[3]. The class distributions follow Gaussian distributions or p-variate heavy-tailed t_ν-distributions with $\nu = 2$ degrees of freedom.

Both the Proposal 1 and Proposal 2 can be used to estimate the regularized class scatter matrices $\hat{\mathbf{\Sigma}}_k(\beta)$ that are needed in RDA rule. We use notation Prop1(ρ, d) and Prop2(ρ, d), where ρ refers to the used loss function and d to the used distance function. To identify the used loss function ρ, we use letters G, H, and T to refer to Gaussian loss ρ_G, Huber's loss ρ_{H} in (13) and Tyler's loss function ρ_{T}, respectively. Furthermore, E and KL indicate that ellipticity distance d_{E} and KL-distance d_{KL}, respectively, are chosen as the distance function d. With the above notation, Prop(G,KL) then refers to original RDA rule based on $\mathbf{S}_k(\beta)$ in (11) and Prop2(T,E), for example, indicates that Tyler's loss function ρ_{T} and ellipticity distance d_{E} are used when estimating the scatter matrices using Proposal 2. For Huber's loss function we used $c^2 = F^{-1}_{\chi^2_p}(0.9)$ as the tuning threshold c.

We compute the estimated misclassification risk as follows. The sample lengths follow multinomial distribution $(n_1, \ldots, n_K) \sim \mathrm{Multin}(N, \mathbf{p})$, where the class probabilities are $p_1 = p_2 = 1/4$ and $p_3 = 1/2$ when $K = 3$ and $p_1 = p_2 = p_3 = 1/6$, and $p_4 = p_5 = 1/4$ when $K = 5$. Then random vectors were drawn from the appropriate class distributions. Each such *training data* set was used to construct the estimated discriminant rules. An additional *test data* set of same sample lengths n_i-s as the training data was generated and classified with the discriminant rules derived from the training set, thereby yielding an estimate of the misclassification risk. For RDA we report the misclassification risk based on the best value of shrinkage parameter β. For each MC trial we compute the

[3]Using different loss functions for different clusters can be advisable when some *a priori* information is available about the class distributions.

Table 1: Average (%) test misclassification errors for unequal spherical covariance matrices ($\boldsymbol{\Sigma}_k = k\mathbf{I}$) for Gaussian (upper table) and t_2-distributed (lower table) clusters. The quantities in subscript inside the parantheses are the standard deviations.

	$K = 3$			$K = 5$		
method	$p = 10$	$p = 20$	$p = 30$	$p = 10$	$p = 20$	$p = 30$
Oracle1	$8.8_{(2.6)}$	$6.2_{(2.3)}$	$4.6_{(1.9)}$	$9.4_{(2.8)}$	$7.7_{(2.8)}$	$6.2_{(2.3)}$
Oracle2	$9.8_{(3.1)}$	$7.6_{(2.6)}$	$6.0_{(2.3)}$	$11.3_{(2.9)}$	$10.1_{(3.3)}$	$9.2_{(2.9)}$
QDA	$19.9_{(4.4)}$	–	–	–	–	–
LDA	$17.1_{(3.8)}$	$20.5_{(4.3)}$	$24.0_{(4.9)}$	$18.8_{(3.8)}$	$24.2_{(4.7)}$	$29.0_{(5.0)}$
Prop1(G,KL)	$12.2_{(3.1)}$	$14.6_{(3.5)}$	$17.9_{(4.3)}$	$15.4_{(3.4)}$	$20.5_{(4.1)}$	$25.8_{(4.8)}$
Prop1(H,KL)	$12.4_{(3.2)}$	$14.6_{(3.5)}$	$17.7_{(4.1)}$	$15.4_{(3.3)}$	$20.3_{(4.1)}$	$25.5_{(4.8)}$
Prop1(T,E)	$10.9_{(3.1)}$	$12.1_{(3.3)}$	$16.5_{(3.9)}$	$13.5_{(3.4)}$	$17.1_{(4.3)}$	$23.9_{(5.1)}$
Prop2(G,E)	$10.5_{(3.0)}$	$11.5_{(3.3)}$	$15.9_{(3.8)}$	$12.9_{(3.4)}$	$16.5_{(4.0)}$	$22.7_{(4.8)}$
Prop2(T,E)	$10.9_{(3.1)}$	$12.1_{(3.3)}$	$16.5_{(3.9)}$	$13.5_{(3.4)}$	$17.1_{(4.3)}$	$23.9_{(5.1)}$
Prop2(H,E)	$10.5_{(3.0)}$	$11.6_{(3.3)}$	$15.7_{(3.8)}$	$12.9_{(3.3)}$	$16.5_{(4.1)}$	$22.6_{(4.8)}$
Prop2(H,KL)	$12.3_{(3.2)}$	$14.8_{(3.6)}$	$18.0_{(4.1)}$	$15.2_{(3.4)}$	$20.1_{(4.2)}$	$25.4_{(4.7)}$
Oracle1	$15.7_{(3.8)}$	$18.2_{(3.9)}$	$21.1_{(4.0)}$	$20.8_{(4.1)}$	$24.5_{(4.2)}$	$27.8_{(4.6)}$
Oracle2	$16.2_{(3.5)}$	$19.1_{(4.2)}$	$21.9_{(4.1)}$	$21.7_{(4.3)}$	$25.8_{(4.1)}$	$29.1_{(4.6)}$
QDA	$26.9_{(5.2)}$	–	–	–	–	–
LDA	$21.8_{(4.9)}$	$25.3_{(5.3)}$	$27.7_{(5.3)}$	$28.6_{(5.6)}$	$32.9_{(5.6)}$	$36.2_{(5.4)}$
Prop1(G,KL)	$19.7_{(4.8)}$	$22.7_{(5.2)}$	$24.7_{(5.1)}$	$27.2_{(5.7)}$	$31.0_{(5.3)}$	$33.8_{(5.4)}$
Prop1(H,KL)	$15.5_{(3.7)}$	$17.9_{(4.0)}$	$20.3_{(4.1)}$	$21.0_{(4.1)}$	$24.6_{(4.5)}$	$28.2_{(4.6)}$
Prop1(T,E)	$16.8_{(4.0)}$	$20.4_{(4.3)}$	$23.4_{(4.7)}$	$23.7_{(4.5)}$	$29.6_{(5.0)}$	$34.0_{(5.3)}$
Prop2(G,E)	$22.3_{(5.9)}$	$24.3_{(5.1)}$	$25.9_{(4.8)}$	$28.1_{(5.4)}$	$32.5_{(5.4)}$	$35.4_{(5.1)}$
Prop2(T,E)	$16.8_{(4.0)}$	$20.4_{(4.4)}$	$23.5_{(4.8)}$	$23.7_{(4.5)}$	$29.7_{(5.0)}$	$34.1_{(5.3)}$
Prop2(H,E)	$16.6_{(3.9)}$	$20.2_{(4.4)}$	$23.6_{(4.6)}$	$23.1_{(4.5)}$	$29.1_{(4.7)}$	$33.8_{(5.3)}$
Prop2(H,KL)	$15.5_{(3.7)}$	$17.9_{(4.0)}$	$20.5_{(4.1)}$	$21.0_{(4.1)}$	$24.6_{(4.4)}$	$28.2_{(4.5)}$

RDA rule (37) for β in the grid $[B] = [0.01, 0.03, \ldots, 0.49, 0.55, 0.60, \ldots, 0.9]$ and the respective estimated misclassification risk. The best value $\beta_0 \in [B]$ is chosen for each RDA approach as the smallest value in the grid $[B]$ that produced the minimum misclassification risk. Reported results are averages over 300 MC trials.

We compare the performance of RDA approaches to conventional LDA and QDA rules as well as to Oracle estimators. We use notation Oracle1 to refer to QDA rule in (36) that uses both the true mean vectors $\boldsymbol{\mu}_k$ and true scatter matrices $\boldsymbol{\Sigma}_k$. Oracle2 denotes QDA rule that uses the true scatter matrices $\boldsymbol{\Sigma}_k$, but estimated mean vectors $\hat{\boldsymbol{\mu}}_k$. For Gaussian samples, $\hat{\boldsymbol{\mu}}_k$ used in Oracle2 are the sample mean vectors and for t_2-distributed samples, they $\hat{\boldsymbol{\mu}}_k$ are the spatial medians of the samples.

7.2. Simulation results

We consider the case of unequal spherical covariance matrices, where the scatter matrix for the kth class is $\boldsymbol{\Sigma}_k = k\mathbf{I}$ for $k = 1, \ldots, K$. This setting is thus somewhat more favourable to QDA, but due to small sample sizes, the performance of QDA does not exceed that of LDA as is shown in Table 1, which

summarizes the simulation results for both the Gaussian and t_2 distributions of the classes. The symmetry center $\boldsymbol{\mu}_1$ of the first class was the origin and for the remaining classes $\boldsymbol{\mu}_k$ were taken to have norm equal to $\delta_k = \|\boldsymbol{\mu}_k\| = 3+k$ in orthogonal directions for Gaussian classes and $\delta_k = \|\boldsymbol{\mu}_k\| = 4+k$ for t_2-distributed classes ($k = 2, \ldots, K$). In the Gaussian case, Prop2(G,E) and Prop2(H,E) are offering consistently the best performance, also outperforming Friedman's original RDA rule, Prop1(G,KL). For example, when $p = 20$, Prop2(G,E) offers 3% improvement in error rate compared to Prop1(G,KL). This illustrates the benefits of choosing the correct penalty (and hence the estimate of joint center covariance matrix): Prop1(G,KL) and Prop2(G,E) are both using the optimal Gaussian loss function, but can have 4% (e.g., case $p = 10$, $K = 5$) difference in the misclassification rate in favor of Prop2(G,E. For Gaussian class distributions, the scale invariant penalty d_{E} offers the best performance. In t_2-case, the results illustrate that robust RDA approaches provide significantly better misclassification rates compared to non-robust RDA approaches using the Gaussian loss function. For example, the best performing robust RDA rules, Prop1(H,KL) and Prop2(H,KL) offer consistently 4–6% improvements to Friedman's Prop1(G,KL). It is somewhat suprising that for t_2-distributed samples, KL-distance is generally performing better than the ellipticity distance. Among RDA approaches, Prop2(G,E) has the worst performance when the class distributions follow the heavy-tailed t_2-distribution.

We then consider the case of equal spherical covariance matrices $\boldsymbol{\Sigma}_k = \mathbf{I}$ for $k = 1, \ldots, K$. In this case, one expects that KL-distance is better choise over ellipticity distance. This set-up favors LDA over QDA due to equality of covariance matrices. The true symmetry center $\boldsymbol{\mu}_1$ of the first class was the origin and for the remaining classes the mean vector $\boldsymbol{\mu}_k$ were taken to have norm equal to δ in orthogonal directions. For Gaussian class distributions, we set $\delta = 3$ and for heavy-tailed t_2-distributions, we set $\delta = 4$. Table 2 gives the estimated misclassification risk for both class distributions. When the class distributions are standard normal distributions, all RDA approaches provide uniformly lower misclassification errors than LDA/QDA, but now the differences between all RDA approaches are insignificant so it is not possible to declare a winner. In general, one can say that all RDA approaches are performing equally well. In t_2-case, the numbers illustrate that robust RDA approaches that are based on KL-distance provide consistently significantly better misclassification risks (about 2–5% improvements) to Prop1(G,KL). Prop2(H,E) is not offering better performance than Prop1(G,KL) despite the robustness of the used loss function. This again illustrates the importance of choosing the right penalty: for equal class covariance matrices (and heavy-tailed distributions), d_{KL} penalty seems more appropriate choice than d_{E}. This observation is also supported by comparing the performance of Prop2(G,E) to Prop1(G,KL) which both are based on Gaussian loss function, but different distance function. Among robust RDA approaches, Prop1(H,I) and Prop2(H,I) are performing the best. It should be noted that for $p = 10$, they offer Oracle performance as their error rates are close to Oracle2 rule which uses the true covariance matrices.

Table 2: Average (%) test misclassification errors for identical spherical ($\mathbf{\Sigma}_k = \mathbf{I}$) covariance matrices in Gaussian (upper table) and t_2-distributed (lower table) samples. The quantities in subscript inside the parantheses are the standard deviations.

	$K = 3$			$K = 5$		
method	$p = 10$	$p = 20$	$p = 30$	$p = 10$	$p = 20$	$p = 30$
Oracle1	$8.9_{(2.9)}$	$9.2_{(3.1)}$	$8.8_{(2.8)}$	$10.9_{(3.2)}$	$10.9_{(3.2)}$	$10.9_{(3.0)}$
Oracle2	$9.9_{(3.1)}$	$10.9_{(3.3)}$	$11.2_{(3.2)}$	$12.9_{(3.2)}$	$14.3_{(3.9)}$	$15.5_{(3.6)}$
QDA	$18.1_{(4.1)}$	–	–	–	–	–
LDA	$11.3_{(3.0)}$	$14.1_{(3.8)}$	$16.9_{(4.2)}$	$14.6_{(3.6)}$	$18.5_{(4.3)}$	$22.9_{(4.7)}$
Prop1(G,KL)	$10.3_{(2.9)}$	$13.0_{(3.6)}$	$15.4_{(3.9)}$	$13.6_{(3.4)}$	$17.4_{(4.1)}$	$21.7_{(4.5)}$
Prop1(H,I)	$10.4_{(3.0)}$	$13.0_{(3.6)}$	$15.4_{(4.0)}$	$13.7_{(3.4)}$	$17.5_{(4.1)}$	$21.8_{(4.5)}$
Prop1(T,E)	$10.8_{(3.1)}$	$13.4_{(3.7)}$	$15.7_{(3.8)}$	$14.5_{(3.4)}$	$18.3_{(4.4)}$	$22.9_{(4.9)}$
Prop2(G,E)	$10.3_{(3.0)}$	$12.9_{(3.7)}$	$15.3_{(3.7)}$	$13.9_{(3.4)}$	$17.6_{(4.1)}$	$22.1_{(4.7)}$
Prop2(T,E)	$10.9_{(3.1)}$	$13.4_{(3.7)}$	$15.6_{(3.8)}$	$14.5_{(3.5)}$	$18.2_{(4.4)}$	$22.9_{(4.8)}$
Prop2(H,E)	$10.4_{(3.0)}$	$13.0_{(3.7)}$	$15.3_{(3.8)}$	$14.0_{(3.4)}$	$17.7_{(4.2)}$	$22.1_{(4.8)}$
Prop2(H,I)	$10.4_{(3.0)}$	$13.0_{(3.6)}$	$15.6_{(4.0)}$	$13.7_{(3.4)}$	$17.5_{(4.1)}$	$21.8_{(4.6)}$
Oracle1	$12.3_{(3.2)}$	$12.5_{(3.5)}$	$12.0_{(3.2)}$	$15.4_{(3.5)}$	$15.4_{(3.4)}$	$15.5_{(3.3)}$
Oracle2	$12.7_{(3.1)}$	$13.3_{(3.6)}$	$13.5_{(3.4)}$	$16.5_{(3.7)}$	$17.3_{(3.5)}$	$18.1_{(3.8)}$
QDA	$24.4_{(5.1)}$	–	–	–	–	–
LDA	$16.2_{(4.1)}$	$19.2_{(4.5)}$	$21.0_{(4.7)}$	$22.1_{(4.7)}$	$25.2_{(5.0)}$	$28.1_{(4.9)}$
Prop1(G,KL)	$14.7_{(4.0)}$	$17.6_{(4.4)}$	$19.2_{(4.3)}$	$20.7_{(4.5)}$	$23.7_{(4.6)}$	$26.6_{(4.8)}$
Prop1(H,KL)	$12.7_{(3.3)}$	$14.6_{(3.8)}$	$16.7_{(3.8)}$	$16.9_{(3.8)}$	$19.5_{(3.7)}$	$22.7_{(4.1)}$
Prop1(T,E)	$14.9_{(4.2)}$	$18.2_{(5.2)}$	$21.0_{(5.7)}$	$20.8_{(4.7)}$	$26.0_{(5.3)}$	$31.0_{(6.1)}$
Prop2(G,E)	$17.5_{(4.9)}$	$20.1_{(5.4)}$	$22.5_{(5.0)}$	$23.6_{(5.2)}$	$28.4_{(5.5)}$	$32.1_{(5.6)}$
Prop2(T,E)	$14.9_{(4.2)}$	$18.3_{(5.2)}$	$21.1_{(5.7)}$	$20.8_{(4.7)}$	$26.1_{(5.3)}$	$31.1_{(6.2)}$
Prop2(H,E)	$14.5_{(4.0)}$	$17.5_{(4.9)}$	$20.4_{(5.2)}$	$19.9_{(4.5)}$	$24.9_{(4.8)}$	$29.5_{(6.1)}$
Prop2(H,KL)	$12.7_{(3.3)}$	$14.7_{(3.7)}$	$16.8_{(3.8)}$	$16.9_{(3.8)}$	$19.5_{(3.7)}$	$22.7_{(4.1)}$

7.3. Data example

For illustrative purposes, we enclose the paper with a simple example of applying RDA on thw well-known Fisher's IRIS data [35] which has $K = 3$ samples, each having $n_k = 50$ $p = 4$-variate observations. We partition the original (3×50) dataset into a training $(3 \times T)$ and a validation $(3 \times V)$ subsets $(T + V = 50)$. The different T/V paritionings used were 30/20, 25/25, 15/35 and 10/40. To demonstrate the robustness of our techniques over the standard Gaussian tools, we replaced two measurements in each training group by outliers with relatively high random amplitudes generated as $\zeta(1, 1, 1, 1)^\top$, where ζ was generated from $Unif(0, 1024)$ for each random T/V splits of the datasets. The training data set is used to estimate the regularized class covariance matrices and forming the RDA rule using 5-fold CV procedure for penalty parameter selection. We then calculated the average misclassification errors on the validation subset and the results, collected in Table 3, are averaged over 100 random T/V partitions of the original dataset. Also results using LDA and QDA rules are reported. These figures clearly illustrate that robust RDA rules outperform the conventional LDA and QDA rules as well as Friedman's RDA rule, Prop1(G,KL). Furthermore, note that RDA rules based on Proposition 2

Table 3: Average (%) validation misclassification errors for IRIS data. Here 30/20, 25/25, 15/35 and 10/40 refer to T/V random splits of $n_k = 50$ measurements in each class into the training and validation subsets. The training data was used for group covariance estimation and forming the RDA rules using 5-fold cross validation. Results are averages of 100 random T/V splits of the data sets.

	30/20	25/25	15/35	10/40
LDA	7.0	6.8	9.6	11.5
QDA	5.0	4.7	6.3	8.3
Prop1(G,KL)	5.1	4.9	7.1	9.3
Prop1(T,E)	2.7	3.6	3.9	4.0
Prop1(H,E)	2.9	3.1	3.9	6.3
Prop1(H,KL)	2.8	3.3	3.9	6.4
Prop2(T,E)	2.8	3.5	3.7	5.8
Prop2(H,E)	2.8	3.1	3.7	4.7
Prop2(H,KL)	2.9	3.4	3.7	5.8

are giving slightly better results compared with RDA rules based on Proposition 1. This is most evident in the case of 10/40 partitioning, which is also the case in which regularization approaches are most useful due to relatively small sample size ($n_k = 10$). For 10/40 partitioning case, Prop2(H,KL) gives 4.2% error rate whereas Prop1(H,KL) attains 6.3% error rate. In constrast, the conventional non-robust LDA and Friedman's Prop1(G,KL) yield 11.5% and 9.3% error rates, respectively.

8. Conclusions

In this paper, we have formulated a joint penalized ML (or M) estimation approach for estimating the unknown scatter matrices of $K > 1$ samples and a joint center. The penalty function is based on a distance that enforces similarity. We considered three different jointly g-convex penalties, namely Riemannian, Ellipticity, and KL-distance in our formulations.

We illustrated the usefullness of our estimators in RDA setting. In this connection, we would like to stress that discriminant analysis is only one application where the developed approach can be used. We expect that our approach and framework can find uses in many other applications such as radar signal processing or graphical models, where similar ideas has been used; See [36, 37], for example. There are still room for improvements in the RDA approach. For example, we did not explore using different loss functions for different classes or using different penalties for different classes. Also, the distance (penalty) function can be different for each class. Such choices can be useful in some applications.

We did not use Frobenius distance which is based on classical Euclidean geometry where as our approach is based on g-convexity which treats $\mathcal{S}(p)$ as a differentiable Riemannian manifold with geodesic path (16). Let us point out that there are other distance functions $d(\mathbf{A}, \mathbf{B})$ that coud be used such as

S-divergence [38]:

$$d_{\mathrm{S}}(\mathbf{A},\mathbf{B}) = \log\Big|\frac{\mathbf{A}+\mathbf{B}}{2}\Big| - \frac{1}{2}\log|\mathbf{A}\mathbf{B}|.$$

S-divergence obviously satistifies D1 and it was shown in [38] that d_{S} is jointly g-convex, i.e., verifies D2. Moreover, S-divergence possesses properties similar to that of geodesic distance $d_{\mathrm{R}}(\cdot,\cdot)$, such as symmetry property D3 (and also affine invariance D4), but has the benefit of being easier to compute. Indeed the induced mean (10) is a solution to a fixed point equation

$$\boldsymbol{\Sigma} = \left(\sum_{k=1}^{K}\pi_k\left(\frac{\boldsymbol{\Sigma}+\boldsymbol{\Sigma}_k}{2}\right)^{-1}\right)^{-1}$$

and thus can be interpreted as weighted harmonic mean of pairwise averages. Despite of the above representation for the mean for fixed $\boldsymbol{\Sigma}_1,\ldots,\boldsymbol{\Sigma}_K$, joint estimation of the scatter matrices result into rather complex estimating equations. Therefore we omitted the use of this distance function in our framework.

Appendix A. Proofs

Proof of Proposition 1 Properties D1, D4 and D5 are obvious. For D2, we show in Lemma 1 below that $\log\mathrm{Tr}(\mathbf{A}^{-1}\mathbf{B})$ is jointly g-convex. Next, we note that $\log|\mathbf{A}^{-1}\mathbf{B}| = \log|\mathbf{A}^{-1}| + \log|\mathbf{B}|$, and that the log-determinant function is a g-linear function, i.e. $\pm\log|\mathbf{A}|$ is g-convex. Hence $\log|\mathbf{A}^{-1}\mathbf{B}|$ is jointly g-convex, and so D2 holds. Since (10) is a sum of g-convex functions, the necessary and sufficient condition for $\boldsymbol{\Sigma}$ to be the solution to (10) is the vanishing of the gradient, $\boldsymbol{\nabla}_{\boldsymbol{\Sigma}}\sum_{i=1}^{K}\pi_k\, d_{\mathrm{E}}(\boldsymbol{\Sigma}_k,\boldsymbol{\Sigma}) = \mathbf{0}$, the solution of which is easily found to be (20).

Lemma 1. $\log\mathrm{Tr}(\mathbf{A}^{-1}\mathbf{B})$ is a jointly strictly g-convex function.

Proof. The geodesic curves connecting $\mathbf{A}_0$ with $\mathbf{A}_1$ and $\mathbf{B}_0$ with $\mathbf{B}_1$ on the Riemannian PSD manifold are given by:

$$\begin{aligned}\mathbf{A}_t &= \mathbf{A}_0^{\frac{1}{2}}\left(\mathbf{A}_0^{-\frac{1}{2}}\mathbf{A}_1\mathbf{A}_0^{-\frac{1}{2}}\right)^t\mathbf{A}_0^{\frac{1}{2}} = \mathbf{A}_0^{\frac{1}{2}}\mathbf{U}_A\mathbf{D}_A^t\mathbf{U}_A^T\mathbf{A}_0^{\frac{1}{2}},\\ \mathbf{B}_r &= \mathbf{B}_0^{\frac{1}{2}}\left(\mathbf{B}_0^{-\frac{1}{2}}\mathbf{B}_1\mathbf{A}_0^{-\frac{1}{2}}\right)^r\mathbf{B}_0^{\frac{1}{2}} = \mathbf{B}_0^{\frac{1}{2}}\mathbf{U}_B\mathbf{D}_B^r\mathbf{U}_B^T\mathbf{B}_0^{\frac{1}{2}},\end{aligned}$$

where the right hand sides are obtain from using the eigenvalue decompositions

$$\mathbf{A}_0^{-\frac{1}{2}}\mathbf{A}_1\mathbf{A}_0^{-\frac{1}{2}} = \mathbf{U}_A\mathbf{D}_A\mathbf{U}_A^T, \quad \text{and} \quad \mathbf{B}_0^{-\frac{1}{2}}\mathbf{B}_1\mathbf{B}_0^{-\frac{1}{2}} = \mathbf{U}_B\mathbf{D}_B\mathbf{U}_B^T.$$

This gives

$$
\begin{aligned}
\log \mathrm{Tr}\mathbf{A}_t^{-1}\mathbf{B}_r &= \log \mathrm{Tr}\mathbf{A}_0^{-\frac{1}{2}}\mathbf{U}_A\mathbf{D}_A^{-t}\mathbf{U}_A^T\mathbf{A}_0^{-\frac{1}{2}}\mathbf{B}_0^{\frac{1}{2}}\mathbf{U}_B\mathbf{D}_B^r\mathbf{U}_B^T\mathbf{B}_0^{\frac{1}{2}} \\
&= \log \mathrm{Tr}\mathbf{A}_0^{-\frac{1}{2}}\mathbf{U}_A\mathbf{D}_A^{-\frac{t}{2}}\mathbf{D}_A^{-\frac{t}{2}}\mathbf{U}_A^T\mathbf{A}_0^{-\frac{1}{2}}\mathbf{B}_0^{\frac{1}{2}}\mathbf{U}_B\mathbf{D}_B^{\frac{r}{2}}\mathbf{D}_B^{\frac{r}{2}}\mathbf{U}_B^T\mathbf{B}_0^{\frac{1}{2}} \\
&= \log \mathrm{Tr}\mathbf{D}_A^{-\frac{t}{2}}\mathbf{U}_A^T\mathbf{A}_0^{-\frac{1}{2}}\mathbf{B}_0^{\frac{1}{2}}\mathbf{U}_B\mathbf{D}_B^{\frac{r}{2}}\cdot\mathbf{D}_B^{\frac{r}{2}}\mathbf{U}_B^T\mathbf{B}_0^{\frac{1}{2}}\mathbf{A}_0^{-\frac{1}{2}}\mathbf{U}_A\mathbf{D}_A^{-\frac{t}{2}} \\
&= \log \mathrm{Tr}\mathbf{C}\mathbf{C}^T, \text{ where } \mathbf{C} = \mathbf{D}_A^{-\frac{t}{2}}\mathbf{U}_A^T\mathbf{A}_0^{-\frac{1}{2}}\mathbf{B}_0^{\frac{1}{2}}\mathbf{U}_B\mathbf{D}_B^{\frac{r}{2}} \\
&= \log \sum_{i,j}\mathbf{C}_{i,j}^2 = \log \sum_{i,j}\left(\mathbf{U}_A^T\mathbf{A}_0^{-\frac{1}{2}}\mathbf{B}_0^{\frac{1}{2}}\mathbf{U}_B\right)_{i,j}^2 (\mathbf{D}_A)_{ii}^{-t}(\mathbf{D}_B)_{jj}^r \\
&= \log \sum_{i,j}\left(\mathbf{U}_A^T\mathbf{A}_0^{-\frac{1}{2}}\mathbf{B}_0^{\frac{1}{2}}\mathbf{U}_B\right)_{i,j}^2 e^{-t\log\left((\mathbf{D}_A)_{ii}\right)+r\log\left((\mathbf{D}_B)_{jj}\right)}.
\end{aligned}
$$

Since the log-sum-exp expression is strictly convex in (t, r), the log-trace function is jointly stricly g-convex. □

Proof of Proposition 2 KL-divergence satisfies D1 and D4. By Lemma 1, $\mathrm{Tr}(\mathbf{A}^{-1}\mathbf{B})$ is jointly strictly g-convex, which implies $d_{\mathrm{KL}}(\mathbf{A}, \mathbf{B})$ is jointly strictly g-convex, i.e., D2 holds. Since (10) is a sum of strictly g-convex functions, the unique minimizer is found by solving $\boldsymbol{\nabla}_{\boldsymbol{\Sigma}}\sum_{i=1}^K \pi_k\, d_{\mathrm{KL}}(\boldsymbol{\Sigma}_k, \boldsymbol{\Sigma}) = \mathbf{0}$, which gives (21).

References

References

[1] J. H. Friedman, Regularized discriminant analysis, J. Amer. Stat. Assoc. 84 (405) (1989) 165–175.

[2] W. Lee, Y. Liu, Joint estimation of multiple precision matrices with common structures, J. Mach. Learn. Res. 16 (2015) 1035–1062.

[3] P. J. Huber, Robust statistics, Wiley, 1981.

[4] R. A. Maronna, Robust M-estimators of multivariate location and scatter, Ann. Stat. 5 (1) (1976) 51–67.

[5] D. E. Tyler, A distribution-free M-estimator of multivariate scatter, Ann. Stat. 15 (1) (1987) 234–251.

[6] P. J. Rousseeuw, Mathematical Statistics and Applications, Reidel, Dordrecht, 1985, Ch. Multivariate Estimation With High Breakdown Point.

[7] P. L. Davies, Asymptotic behaviour of S-estimates of multivariate location parameters and dispersion matrices, Ann. Stat. (1987) 1269–1292.

[8] K. S. Tatsuoka, D. E. Tyler, On the uniqueness of S-functionals and M-functionals under nonelliptical distributions, Ann. Stat. (2000) 1219–1243.

[9] A. Wiesel, Geodesic convexity and covariance estimation, IEEE Trans. Signal Process. 60 (12) (2012) 6182–6189.

[10] A. Wiesel, Unified framework to regularized covariance estimation in scaled gaussian models, IEEE Trans. Signal Process. 60 (1) (2012) 29–38.

[11] T. Zhang, A. Wiesel, M. S. Greco, Multivariate generalized Gaussian distribution: Convexity and graphical models, IEEE Trans. Signal Process. 61 (16) (2013) 4141–4148.

[12] S. Sra, R. Hosseini, Conic geometric optimization on the manifold of positive definite matrices, SIAM Journal on Optimization 25 (1) (2015) 713–739.

[13] E. Ollila, D. E. Tyler, Regularized M-estimators of scatter matrix, IEEE Trans. Signal Process. 62 (22) (2014) 6059–6070.

[14] A. Wiesel, T. Zhang, Structured robust covariance estimation, Foundations and Trends in Signal Processing 8 (3) (2015) 127–216.

[15] G. Frahm, Generalized elliptical distributions: theory and applications, Ph.D. thesis, Universität zu Köln (2004).

[16] E. Ollila, D. Tyler, V. Koivunen, H. Poor, Complex elliptically symmetric distributions: survey, new results and applications, IEEE Trans. Signal Process. 60 (11) (2012) 5597–5625.

[17] F. R. Hampel, E. M. Ronchetti, P. J. Rousseeuw, W. A. Stahel, Robust Statistics: The Approach Based on Influence Functions, Wiley, New York, 1986.

[18] R. A. Maronna, R. D. Martin, V. J. Yohai, Robust Statistics: Theory and Methods, Wiley, New York, 2006.

[19] J. T. Kent, D. E. Tyler, Redescending M-estimates of multivariate location and scatter, Ann. Stat. 19 (4) (1991) 2102–2119.

[20] P. J. Huber, Robust estimation of a location parameter, Ann. Math. Stat. 35 (1) (1964) 73–101.

[21] D. E. Tyler, Statistical analysis for the angular central Gaussian distribution on the sphere, Biometrika 74 (3) (1987) 579–589.

[22] P.-A. Absil, R. Mahony, R. Sepulchre, Optimization algorithms on matrix manifolds, Princeton University Press, 2009.

[23] R. Bhatia, Positive definite matrices, Princeton University Press, 2009.

[24] M. Moakher, A differential geometric approach to the geometric mean of symmetric positive-definite matrices, SIAM J. Matrix Anal. Appl. 26 (3) (2005) 735–747.

[25] N. J. Higham, Functions of matrices: theory and computation, SIAM, 2008.

[26] D. A. Bini, B. Iannazzo, Computing the Karcher mean of symmetric positive definite matrices, Linear Algebra Appl. 438 (4) (2013) 1700–1710.

[27] T. Zhang, A majorization-minimization algorithm for the Karcher mean of positive definite matrices, arXiv preprint arXiv:1312.4654 (2013).

[28] T. M. Cover, J. A. Thomas, Elements of information theory, John Wiley & Sons.

[29] W. James, C. Stein, Estimation with quadratic loss, Proceedings of the fourth Berkeley symposium on mathematical statistics and probability 1 (1961) 361–379.

[30] Y. Sun, P. Babu, D. P. Palomar, Regularized Tyler's scatter estimator: Existence, uniqueness, and algorithms, IEEE Trans. Signal Process. 62 (19) (2014) 5143–5156.

[31] D. R. Hunter, K. Lange, A tutorial on mm algorithms, Am. Stat. 58 (1) (2004) 30–37.

[32] M. Razaviyayn, M. Hong, Z.-Q. Luo, A unified convergence analysis of block successive minimization methods for nonsmooth optimization, SIAM J. Optim. 23 (2) (2013) 1126–1153.

[33] L. Dümbgen, K. Nordhausen, H. Schuhmacher, New algorithms for M-estimation of multivariate scatter and location, J. Mult. Anal. 144 (2016) 200–217.

[34] B. M. Brown, Statistical uses of the spatial median, J. Royal Stat. Soc., Ser. B (1983) 25–30.

[35] R. A. Fisher, The use of multiple measurements in taxonomic problems, Annals of eugenics 7 (2) (1936) 179–188.

[36] O. Besson, S. Bidon, J.-Y. Tourneret, Covariance matrix estimation with heterogeneous samples, IEEE Trans. Signal Process. 56 (3) (2008) 909–920.

[37] P. Danaher, P. Wang, D. M. Witten, The joint graphical lasso for inverse covariance estimation across multiple classes, J. Royal Stat. Soc., Ser. B 76 (2) (2014) 373–397.

[38] S. Sra, Positive definite matrices and the S-divergence, arXiv preprint arXiv:1110.1773.